

REMARKS

Claims 1-34 remain pending in the application. Reconsideration is respectfully requested in light of the following remarks.

Section 103(a) Rejection:

The Examiner rejected claims 1-7, 10, 11, 14, 15, 18-24 and 27-32 under 35 U.S.C. § 103(a) as being unpatentable over Doyle (U.S. Patent 6,009,455) in view of IEEE Conference Proceeding “When Peer-to-Peer comes Face-to-Face: Collaborative Peer-to-Peer Computing in Mobile Ad Hoc Networks”, by Kortuem et al., published in First International Conference on Peer-to-Peer Computing proceedings, August 2001, Pages 75-91 (hereinafter “IEEE”), and further in view of Kampe, et al. (U.S. Publication 2002/0042693) (hereinafter “Kampe”). Applicants respectfully traverse this rejection for at least the reasons below.

Doyle’s system would be recognized in the art as a **conventional grid computing system** as disclosed in the Background section of the instant application. As further explained below, Doyle (whether considered alone or in combination with the other references) simply describes the same sort of **conventional** grid computing system that is described in the Background section of the present application. The IEEE reference pertains to decentralized ad hoc mobile peer networks. The Kampe reference is directed at “clusters”, which Kampe defines as “high availability computer systems” or “carrier-grade high availability platforms.” Kampe’s cluster consists of peer nodes that cooperate to provide a highly available service. Contrary to the Examiner’s assertion, Kampe’s cluster is not equivalent in structure or purpose to a *grid computing system* as defined in Applicants’ specification and recited in Applicants’ claims. Doyle, IEEE and Kampe, alone or in combination, do not teach or suggest Applicants’ claimed invention, as further shown below.

Regarding claim 1, contrary to the Examiner’s assertion, the cited

references, alone or in combination, do not teach or suggest *a grid comprising one or more compute nodes already configured to participate in the grid; a node, wherein when the node is not configured with the current software to perform one or more compute jobs to participate in the grid as a compute node, the node is configured to discover the master node in accordance with one or more peer-to-peer platform protocols and, in response to said discovering the master node, send information about the node to the discovered master node in accordance with the one or more peer-to-peer platform protocols.*

Applicants will first respond to the Examiner's arguments in regard to Doyle and Kampe, and then provide responses to the Examiner's arguments in regard to the IEEE reference.

The Examiner cites Doyle, reference character 12 in Fig. 2a (as the "client control program"), and Doyle, col. 3, line 64 - col. 4, line 10 as "sending information"). However, the Doyle reference clearly discloses that Doyle's "client computers" (reference character 11 in Fig. 2a) are already configured as "compute nodes" in Doyle's "grid". See, for example, col. 3, lines 50-52. Applicants note that, in the Action dated August 26, 2008, the Examiner admits that Doyle's client computers are "each already configured to participate in the grid with current[ly] executable software to perform one or more compute jobs."

The Examiner asserts Kampe discloses a method of joining a cluster, and equates Kampe's "cluster" to the instant application's "grid", citing paragraphs [0039]-[0044] and Figure 3. (Applicants note that the Examiner appears to intend Figure 2 of Kampe in the Office Action where Figure 3 is cited.) Applicants traverse the assertion that Kampe's "cluster" is equivalent to "grid" as used in the instant application. **The Kampe reference is directed at "clusters"**, which Kampe defines as "high availability computer systems" or "carrier-grade high availability platforms." Kampe's clusters consist of nodes that cooperate to provide highly available services. Contrary to the Examiner's assertion, Kampe's cluster is **not** equivalent in structure or purpose to a *grid* or to a *grid*

computing system as defined in Applicants' specification and recited in Applicants' claims and as the term *grid computing system* is understood in the art.

In the Action dated April 8, 2008, the Examiner asserted that Kampe discloses "a node is not configured to participate in the grid as a compute node, (a node attempts to join a cluster at step 110 in Figure 3 [should be Figure 2], page 3, paragraph [0039])." In the Action dated August 26, 2008, the Examiner asserts that Kampe discloses "A node is not configured with configuration data (an entity joining a cluster obtains its configuration data from the master cluster membership monitoring entity, see, e.g., page 3, paragraph [0044])." In the Response to Arguments section of the latter Office Action, the Examiner further asserts "Kampe teaches discovering a master node and receiving a configuration data from the discovered master node for a node is not fully configured as a cluster member (see, e.g., page 3, paragraph [0040]-0044))."

Contrary to the Examiner's assertions, Kampe clearly discloses that a node that attempts to join a cluster must be pre-configured to participate in the cluster. Kampe's teaching of a method of joining a cluster, as illustrated in Figure 2, and elsewhere in the Kampe reference, clearly describes that nodes are, **and must be**, pre-configured as cluster nodes before joining a cluster. Paragraph [0039] states, as a first step of a method of joining a cluster (emphasis added):

A peer node boots and the local cluster membership monitor entity is started as part of the carrier-grade highly available middleware, step 100.

Referring to Fig. 1 of Kampe, the "CMM Entity", elements 50, 52 and 54, clearly represent the cluster membership monitor (see, e.g., paragraph [0030], first sentence). Thus, paragraph [0039] makes it clear that an instance of the cluster membership monitor must reside on the peer node at startup. Kampe's method of joining a cluster depends on the node being pre-configured with a local instance of Kampe's cluster membership monitor. Further note that the citation states that the local cluster membership monitor is started as part of the carrier-grade highly available middleware. This clearly indicates that the node must be pre-configured with a particular "carrier-grade highly available middleware" (i.e. the middleware must be installed prior to startup) to join as a member

node in Kampe's cluster or "carrier-grade high availability platform." The next sentence of paragraph [0039] states (emphasis added):

The initial default hold-off, time-out, and retry periods are pre-configured in the boot image or NVRAM.

Thus, these parameters used in Kampe's clusters are pre-configured in a node at the start of the method of joining a cluster.

In addition, other components of Kampe's cluster nodes are also clearly described as being pre-configured on Kampe's nodes. For example, paragraph [0029] states (emphasis added):

On each node 20, 22, 24, and 26, the node id 30, node type 32 and master priority 34 are known at cluster membership monitor initialization. This information can be retrieved from Non-Volatile Random Access Memory (NVRAM) or the boot image.

Cluster membership monitor initialization is what is referred to in paragraph [0039], step 100 of Figure 2 ("the local cluster membership monitor entity is started"). Paragraphs [0031]-[0038] also further describe pre-configured nodes. Again, Kampe clearly discloses, in multiple places and consistently throughout the reference, that a node must be pre-configured as a cluster node to join a cluster.

In response to the Examiner's assertion in the Action dated August 26, 2008 that Kampe discloses "A node is not configured with configuration data (an entity joining a cluster obtains its configuration data from the master cluster membership monitoring entity, see, e.g., page 3, paragraph [0044])," and "Kampe teaches discovering a master node and receiving a configuration data from the discovered master node for a node is not fully configured as a cluster member (see, e.g., page 3, paragraph [0040]-0044)," paragraph [0044] of Kampe states:

A cluster membership monitor entity joining a cluster obtains its configuration data from the master cluster membership monitor entity and sets its time-out, retry and consensus protocol times accordingly. These override the default parameters used during initialization.

From Kampe's teachings and from Applicants' above arguments, it is clear that Kampe's cluster membership monitor entity only receives configuration data from the master cluster membership monitor entity; the configuration data simply overrides default parameters used when the cluster membership monitor entity initialized. The cluster membership monitor entity is clearly already configured with any and all necessary software. Applicants again refer to paragraph [0039], which states, as a first step of a method of joining a cluster (emphasis added):

A peer node boots and the local cluster membership monitor entity is started as part of the carrier-grade highly available middleware, step 100.

In direct contrast to Kampe's teachings in the citations provided by the Examiner, Applicants' claim 1 recites a node, wherein when the node is not configured with the current software to perform one or more compute jobs to participate in the grid as a compute node. **Kampe's nodes are clearly pre-configured with "current software."**

The Examiner further asserts that Kampe discloses "in response to said discovering the master node, send information about the node to the discovered master node in accordance with the one or more peer-to-peer platform protocols." The Examiner asserts "the node attempts to join a cluster sends node information...to the master node which maintains a cluster configuration repository," citing page 2, paragraph [0028]. **The citation does not disclose what the Examiner asserts it discloses.** The citation does not disclose that a node attempting to join a cluster sends node information to the master node. The citation instead discloses that "node configuration is maintained in a cluster configuration repository." Paragraph [0044] states that "a node can only join a cluster if its node id is configured in the cluster configuration repository." This strongly indicates that the node information is and must be already present in the cluster configuration repository for the node to join the cluster. A further indication that node information for a node that may attempt to join a cluster is already present in the cluster configuration repository is found in paragraph [0027], which states:

The cluster membership monitor is also provisioned with information about potential cluster members.

Furthermore, in paragraph [0041], Kampe discloses “a search for the master begins.” In paragraph [0042], Kampe goes on to state “If there is consensus, step 130, and a master is located, step 150, the peers form a cluster and the master logs the ‘cluster join’ event to its local system log, step 160.” Kampe does not mention the node sending node information to the master node when describing the method of joining a cluster.

The Examiner further asserts Kampe discloses “the master node is further configured to send grid configuration information to the node,” and goes on to assert “a node attempts to join the cluster obtains its configuration data from the master cluster node,” citing page 3, paragraph [0044]. The cited paragraph states:

A cluster membership monitor entity joining a cluster obtains its configuration data from the master cluster membership monitor entity and sets its time-out, retry and consensus protocol timers accordingly. These override the default parameters used during initialization.

The configuration data obtained from the master node appears to consist of only specific settings for time-out, retry and consensus protocol timers, settings that were pre-configured as described above (paragraph [0039] states “the initial default hold-off, time-out, and retry periods are pre-configured in the boot image or NVRAM.”) The obtained configuration data is simply used to override the default, pre-configured settings. Applicants note that claim 1 goes on to recite:

wherein the node is further configured to, in response to said grid configuration information received from the discovered master node, self-configure as a compute node in the grid in accordance with the grid configuration information, wherein after self-configuring as a compute node the node is configured with the current software to perform one or more compute jobs in the grid.

Kampe clearly does not disclose this limitation as recited in Applicants’ claim. As Applicants have made clear in the above discussion of Kampe’s method of joining a cluster, Kampe’s nodes are clearly described as being pre-configured as cluster nodes prior to joining a cluster, with all necessary components including software being installed on the node prior to startup; even the settings that may be overridden by “configuration data” have pre-configured default values.

The Examiner goes on to assert “it would have been obvious to combine Kampe with Doyle in view of IEEE in order for a node to automatically join the existing grid and configure itself from the discovered master node which maintains grid configuration information.” Applicants note that the “grid configuration information” maintained by the master node is the cluster configuration repository. Paragraph [0028] of Kampe describes the repository:

Node configuration is maintained in a cluster configuration repository, a copy of which is located on the master and vice-master nodes. Within the repository, a cluster 10 is represented as a table of rows, one row per node. In a preferred embodiment each row contains at least the following information:

- node name used for log purposes only
- node id unique identifier associated with a node
- node address communications address of the node (IP address)
- node type non-peer, NSP, HSP
- node priority priority for assuming master role (0 = non master)
- timeout heartbeat timer for this node
- retry count number of missed heartbeats before initiating cluster consensus procedure

Again, paragraph [0044] states “a node can only join a cluster if its node id is configured in the cluster configuration repository.” Paragraph [0027] states, “The cluster membership monitor is also provisioned with information about potential cluster members.” Thus, the “grid configuration information” maintained by the master node (the repository) is simply a list of cluster members and potential cluster members, clearly referring to nodes pre-configured to become cluster members. Note that this is all pre-configured information. **“A node can only join a cluster if its node id is configured in the cluster configuration repository.”** Contrary to Examiner’s assertion, Kampe does not teach that a node “automatically joins an existing grid and configures itself from the discovered master node which maintains grid configuration information.” Kampe instead teaches that only a pre-configured node can join a cluster, and the node can only join the cluster if the cluster itself is pre-configured to indicate that the node can join the cluster. Thus, combining Kampe with Doyle in view of IEEE would not and could not achieve the results the Examiner asserts.

In the Action dated August 26, 2008, the Examiner goes on to assert “Because Doyle taught the one or more computer node[s] already configured with currently executable software to perform compute jobs as presented above[,] it would have been obvious to receive the executable software as the configuration information in order to properly participate as a compute node.” The Examiner’s assertion here makes no sense. Clearly, there would be no need for a node already configured with currently executable software to perform compute jobs, as taught by Doyle, to receive the executable software “as the configuration information” in order to properly participate as a compute node, as asserted by the Examiner.

The Examiner relies on the IEEE reference to teach discovering the master node of the grid. **However, neither Doyle’s, nor IEEE’s, nor Kampe’s teachings (whether considered alone or in combination) pertain to or are suggestive of a grid in which some nodes are already configured to participate in the grid as compute nodes with current software to perform one or more compute jobs, and another node not configured with the current software to perform one or more compute jobs to participate in the grid as a compute node discovers the master node.**

Furthermore, IEEE explicitly teaches away from discovery of a master node according to a peer-to-peer platform protocol. In contrast, IEEE specifically teaches discovery for decentralized mobile peer networks. IEEE states that decentralization “is not a mere option for mobile peer-to-peer networks, but a necessity.” (See last paragraph of section 3.3.4 cited by the Examiner). Thus, one of ordinary skill in the art would not seek to combine the teachings of IEEE with the other references since IEEE explicitly teaches away from master node discovery. Moreover, such a combination would make no sense because the nodes in Doyle are already configured as compute nodes, and the nodes in Kampe are pre-configured as cluster nodes. There would be no reason for a node in Doyle to discover the master node to obtain grid configuration information including current software for self-configuration as a compute node because the nodes are pre-configured as compute nodes, nor is there any reason in Kampe for a node to discover

the master node to obtain cluster configuration information including current software for self-configuration as a cluster node because the nodes are pre-configured as cluster nodes.

Even if the teachings of IEEE were applied to the other references, it would not result in Applicants' claimed invention. IEEE teaches a mobile peer discovering another decentralized mobile peer to form an ad hoc network. At most, applying these teachings to Doyle would result in one compute node discovering another peer compute node, not the master node, and applying these teachings to Kampe would result in one cluster node discovering another peer cluster node, not the master node. Moreover, as discussed above, none of the references have anything to do with a node discovering the master in order to obtain grid configuration to self-configure itself as a compute node in the grid.

Also, the teachings of IEEE pertain to the formation of ad hoc decentralized peer networks. The term "ad hoc" refers to something done for a particular purpose only when the situation makes it necessary or desirable, rather than being arranged in advance or being part of a general plan. The grid in Doyle is not an ad hoc system, nor is the cluster in Kampe an ad-hoc system. Thus, the teachings of IEEE would not pertain to Doyle's system or to Kampe's system. The concept of ad hoc discovery makes no sense in the context of Doyle or Kampe.

Furthermore, the cited art, alone or in combination, does not teach or suggest a node, in response to discovering the master node, sending information about the node to the discovered master node. Doyle discloses, at col. 3, lines 58-61 (emphasis added):

When the client control program is initially activated or is in the idle state, it executes an availability algorithm. The primary function of the availability algorithm is to notify the master computer that the client is available.

Clearly, Doyle does not teach or suggest, in this citation or elsewhere, that the availability algorithm sends information about the client computer 11 to the master computer **in response to discovering the master computer**. To the contrary, Doyle explicitly

teaches that the client control program sends availability (and other) information to the master computer when the client control program is initially activated or is in the idle state. Furthermore, Doyle does not teach or suggest that a client computer 11 would have any need to discover the master computer 5. Doyle discloses that the client computer 11 is pre-configured with a “client control program 12”. The above citation discloses that the availability algorithm on Doyle’s client computer 11 notifies the master computer that the client is available “when the client control program is initially activated or is in the idle state”. Doyle’s client computer 11 is already aware of the master computer 5, and thus it would not be necessary for Doyle’s client computer to perform any discovery operation in accordance with a peer-to-peer platform protocol to discover the master computer.

The cited references, alone or in combination, do not teach or suggest *a node discovering the master node and, in response to said discovering the master node, sending information about the node to the discovered master node*. The Examiner acknowledges that Doyle does not disclose *discovering the master node in accordance with one or more peer-to-peer platform protocols*, and relies on IEEE to disclose “discovery” and “peer-to-peer platform protocols”. However, combining IEEE’s method of discovery with Doyle’s system **would not make sense**. For example, Doyle’s client computers, which include a “client control program” are clearly **pre-configured** as “compute nodes” in Doyle’s grid. Furthermore, the Doyle reference does not teach or suggest that Doyle’s “client computers 11” have any need for discovery of the master computer. Doyle’s pre-configured “client computers 11”, from Doyle’s description, are already aware of the master computer, and thus Doyle’s system would have no need for discovery as disclosed in IEEE.

In further regard to claim 1, contrary to the Examiner’s assertion, the cited references, alone or in combination, do not teach or suggest *wherein the master node is further configured to, in response to said information about the node, send grid configuration information to the node in accordance with the one or more peer-to-peer platform protocols*. The Examiner cites Doyle, “job computation module”, reference

character 14 in Fig. 2c, and asserts “determine which mode the program should operate based on the job request message from job request means, 1 in Fig. 2c, and sends it to the available clients”, citing Doyle, col. 5, line 64 to col. 6, line 16. Contrary to the Examiner’s assertion, Doyle, in the cited portions or elsewhere, does not teach that the master computer 5 sends grid configuration information to the node. Instead, Doyle discloses that the master computer 5 sends “commands and files” to the “selected clients 20” of FIG. 2d. The “selected clients 20” are the same as the “available clients 17” of Fig. 2b and the “client computers 11” illustrated in Fig 2a, or at least a selected subset of the clients 17 and client computers 11. Thus, Doyle’s “selected clients 20” that receive “commands and files” (col. 6, lines 21-24) from the master computer are **already pre-configured** as “compute nodes” in Doyle’s grid. In FIGs 2c and 2d and the accompanying discussion that includes the Examiner’s cited portions, Doyle clearly discloses that the “commands and files” sent to the “selected clients 20” are intended to configure the selected clients 20 to distribute a portion of a particular job to the selected clients 20 for execution, and do not include “grid configuration information” as required in claim 1 of the instant application. This is made clear in col. 6, lines 24-29:

Each selected client is downloaded with the **job request files** 18 included in the segment group package 19 whose contents are **based on the job request signal 40**. FIG. 2d depicts one select client receiving the files constituting segment group package A, while the other selected client receives the files constituting segment group package B.

Again, Doyle’s “selected clients 20” already include the client control program 12 and are already aware of the master computer and thus are already **pre-configured** as “compute nodes” in Doyle’s grid. There is no need for Doyle’s master computer to send grid configuration information to the “selected clients 20” so that the selected nodes can “self-configure” as compute nodes in Doyle’s grid; Doyle’s “selected clients” are already configured as “compute nodes” in Doyle’s grid. FIGs 2c and 2d of Doyle simply illustrate the distribution of portions of jobs to the selected clients 20 in Doyle’s grid.

In further regard to claim 1, contrary to the Examiner’s assertion, the cited references, alone or in combination, does not teach or suggest *wherein the node is*

further configured to, in response to said grid configuration information received from the discovered master node, self-configure as a compute node in the grid in accordance with the grid configuration information, wherein after self-configuring as a compute node the node is configured with the current software to perform one or more compute jobs in the grid.

In regard to the Doyle reference, the Examiner asserts “availability algorithm, 13 in FIG. 2b, concludes the respective client computer as available clients and qualification algorithm, 45 in FIG. 2b, determines the available clients as a candidate to participate in a distributed computation, see e.g., col. 4, lines 11-27”. As noted above, Doyle’s “client computers 11” and “available clients 17” are **pre-configured** as “**compute nodes**” in **Doyle’s grid**. Doyle’s client computers 11 in FIG. 2a **already include Doyle’s client control program 12**. The Examiner’s assertion that Doyle’s system, “[the] availability algorithm concludes the respective client computer as available clients and qualification algorithm...determines the available clients as a candidate to participate in a distributed computation” simply describes Doyle’s system determining a set of “selected clients 20” to perform portions of computations for a job. The Examiner’s assertion has nothing to do with a node that is not configured as a compute node in a grid, in response to grid configuration information sent by a discovered master node in the grid, self-configuring as a compute node in the grid in accordance with the grid configuration information received from the discovered master node, wherein after self-configuring as a compute node the node is configured with the current software to perform one or more compute jobs in the grid.

Applicants also assert that one of ordinary skill in the art would have had no reason to combine the references in the manner proposed by the Examiner. For example, Doyle’s client computers, which include a “client control program,” are clearly **pre-configured** as “compute nodes” in Doyle’s grid. The Doyle reference does not teach or suggest that Doyle’s “client computers 11” have any need for discovery of the master computer. Doyle’s pre-configured “client computers 11”, from Doyle’s description, are already aware of the master computer, and thus Doyle’s system would have no need for

discovery as disclosed in IEEE. Thus, the proposed combination makes no sense. (Regarding the Kampe reference, Kampe already discloses locating other cluster nodes, so there is no reason to combine Kampe with IEEE for discovery purposes.) Also, as shown above, the IEEE reference teaches away from master node discovery and does not pertain to the types of systems in Doyle and Kampe. Moreover, the Examiner's stated reason for combining the Doyle and IEEE references is merely conclusory. And in any case, as noted above, the cited references, alone or in combination, do not teach all of the limitations found in claim 1 of the instant application.

Thus, for at least the reasons presented above, the rejection of claim 1 is not supported by the cited art and removal thereof is respectfully requested.

In regard to claim 10, the cited references, alone or in combination, do not teach or suggest a master node configured to communicate with one or more of a plurality of nodes not configured to participate in the grid as compute nodes in accordance with one or more peer-to-peer platform protocols to configure the one or more of the plurality of nodes to participate as compute nodes in the grid computing system.

In regard to the Kampe reference, Applicants refer to the responses to these assertions given above in regards to claim 1. **Again, contrary to the Examiner's assertion, Kampe clearly discloses that a node that attempts to join a cluster must be pre-configured to participate in the cluster.**

As discussed above in regard to claim 1, the Doyle reference clearly discloses that Doyle's "client computers" (reference character 11 in Fig. 2a) are already configured as "compute nodes" in Doyle's "grid". **Thus, there would be no need in Doyle's system for the master computer to communicate with Doyle's client systems to configure the client systems as "compute nodes" in Doyle's "grid".** Doyle's client systems are described as being pre-configured as "compute nodes" in Doyle's "grid".

Combining IEEE's method of discovery with Doyle's system **would not result in what is recited in claim 10 of the instant application.** For example, Doyle's client computers, which include a "client control program," are clearly **pre-configured** as "compute nodes" in Doyle's grid. The Doyle reference does not teach or suggest that any need for "discovery of all neighbor devices." Doyle's pre-configured "client computers 11", from Doyle's description, are already aware of the master computer, and thus Doyle's system would have no need for discovery as disclosed in IEEE. **Furthermore, even if the Doyle and IEEE references were properly combinable, the combination would not produce anything like what is recited in claim 10 of the present application.**

Applicants also assert that one of ordinary skill in the art would have had no reason to combine the references in the manner proposed by the Examiner. For example, Doyle's client computers, which include a "client control program," are clearly **pre-configured** as "compute nodes" in Doyle's grid. The Doyle reference does not teach or suggest that Doyle's "client computers 11" have any need for discovery of the master computer. Doyle's pre-configured "client computers 11", from Doyle's description, are already aware of the master computer, and thus Doyle's system would have no need for discovery as disclosed in IEEE. Thus, the proposed combination makes no sense. (Regarding the Kampe reference, Kampe already discloses locating other cluster nodes, so there is no reason to combine Kampe with IEEE for discovery purposes.) Also, as shown above, the IEEE reference teaches away from master node discovery and does not pertain to the types of systems in Doyle and Kampe. Moreover, the Examiner's stated reason for combining the Doyle and IEEE references is merely conclusory. And in any case, as noted above, the cited references, alone or in combination, do not teach all of the limitations found in claim 10 of the instant application.

Thus, for at least the reasons presented above, the rejection of claim 10 is not supported by the cited art and removal thereof is respectfully requested.

In regard to claim 14, the cited references, alone or in combination, do not teach or suggest *a node, comprising: a processor; and a memory comprising program instructions, wherein the program instructions are executable by the processor to: when the node is not configured to participate in a grid as a compute node, discover a master node in accordance with one or more peer-to-peer platform protocols, wherein the master node is configured to manage the grid, wherein the grid comprises one or more compute nodes already configured to participate in the grid; and in response to said discovering the master node, send information about the node to the discovered master node in accordance with the one or more peer-to-peer platform protocols.*

In regard to the Examiner's assertions regarding the Kampe reference, Applicants refer to the responses to these assertions given above in regards to claim 1. **Again, contrary to the Examiner's assertion, Kampe clearly discloses that a node that attempts to join a cluster must be pre-configured to participate in the cluster.**

The Examiner acknowledges that Doyle does not disclose *discovering the master node in accordance with one or more peer-to-peer platform protocols*, and relies on IEEE to disclose "discovery" and "peer-to-peer platform protocols".

The Doyle reference clearly discloses that Doyle's "client computers" (reference character 11 in Fig. 2a) are already configured as "compute nodes" in Doyle's "grid", as acknowledged by the Examiner (refer to the response to the rejection of claim 1). Client systems are **pre-configured as "compute nodes" in Doyle's grid system**. The "node" recited in claim 14, at discovery of the master node, is **not configured as a compute node in the grid computing system**.

Furthermore, the cited art, alone or in combination, does not teach or suggest a node, in response to discovering the master node, sending information about the node to the discovered master node. Doyle discloses, at col. 3, lines 58-61 (emphasis added):

When the client control program is initially activated or is in the idle state, it executes an availability algorithm. The primary function of the availability algorithm is to notify the master computer that the client is available.

Doyle does not teach or suggest, in this citation or elsewhere, that the availability algorithm sends information about the client computer 11 to the master computer **in response to discovering the master computer**. Doyle clearly teaches that the client control program sends availability (and other) information to the master computer when the client control program is initially activated or is in the idle state. Furthermore, Doyle does not teach or suggest that a client computer 11 would have any need to discover the master computer 5. Doyle discloses that the client computer 11 is pre-configured with a “client control program 12”. The above citation discloses that the availability algorithm on Doyle’s client computer 11 notifies the master computer that the client is available “when the client control program is initially activated or is in the idle state”. Doyle’s client computer 11 appears to be already aware of the master computer 5, and thus it would not be necessary for Doyle’s client computer to discover the master computer.

The cited references, alone or in combination, do not teach or suggest *a node discovering the master node and, in response to said discovering the master node, sending information about the node to the discovered master node*. The Examiner acknowledges that Doyle does not disclose *discovering the master node in accordance with one or more peer-to-peer platform protocols*, and relies on IEEE to disclose “discovery” and “peer-to-peer platform protocols”. However, combining IEEE’s method of discovery with Doyle’s system **would not result in the above limitation**. For example, Doyle’s client computers, which include a “client control program” are clearly **pre-configured** as “compute nodes” in Doyle’s grid. Furthermore, the Doyle reference does not teach or suggest that Doyle’s “client computers 11” have any need for discovery of the master computer. Doyle’s pre-configured “client computers 11”, from Doyle’s description, are already aware of the master computer, and thus Doyle’s system would have no need for discovery as disclosed in IEEE.

In further regard to claim 14, the cited references, alone or in combination, do not teach or suggest the node *receiving grid configuration information from the master node in accordance with the one or more peer-to-peer platform protocols*. The Examiner asserts “client control program, reference character 12 in Figure 2a, sends to the master computer the existence and configuration of various predetermined resources on the client computer, see e.g. col. 3, line 64 to col. 4, line 10.” The “client control program” sending information to the master computer clearly does not teach or suggest the node receiving grid configuration information from the master node.

Contrary to the Examiner’s assertion, Doyle, in the cited portions or elsewhere, does not teach that the master computer 5 sends grid configuration information to the node. Instead, Doyle discloses that the master computer 5 sends “commands and files” to the “selected clients 20” of FIG. 2d. Doyle’s “selected clients 20” that receive “commands and files” (col. 6, lines 21-24) from the master computer are pre-configured as “compute nodes” in Doyle’s grid. In FIGs 2c and 2d and the accompanying discussion, Doyle clearly discloses that the “commands and files” sent to the “selected clients 20” are intended to configure the selected clients 20 to distribute a portion of a particular job to the selected clients 20 for execution, and do not include “grid configuration information” as disclosed in claim 14 of the instant application. This is made clear in col. 6, lines 24-29 of Doyle.

Again, Doyle’s “selected clients 20” already include the client control program 12 and are already aware of the master computer and thus are pre-configured as “compute nodes” in Doyle’s grid. There is no need for Doyle’s clients to receive grid configuration information from the master node so that the nodes can “self-configure” as compute nodes in Doyle’s grid; Doyle’s “selected clients” are already configured as “compute nodes” in Doyle’s grid.

In further regard to claim 14, the cited references, alone or in combination, does not teach or suggest *wherein the node is further configured to, in response to said grid configuration information, self-configure as a compute node in the grid in*

accordance with the grid configuration information. The Examiner asserts “availability algorithm, 13 in FIG. 2b, concludes the respective client computer as available clients and qualification algorithm, 45 in FIG. 2b, determines the available clients as a candidate to participate in a distributed computation, see e.g., col. 4, lines 11-27”. As noted above, Doyle’s “client computers 11” and “available clients 17” are **pre-configured** as **compute nodes** in Doyle’s grid. Doyle’s client computers 11 in FIG. 2a **already include Doyle’s client control program 12.** The Examiner’s assertion that Doyle’s system, “[the] availability algorithm concludes the respective client computer as available clients and qualification algorithm...determines the available clients as a candidate to participate in a distributed computation” simply describes Doyle’s system determining a set of “selected clients 20” to perform portions of computations for a job. The Examiner’s assertion has nothing to do with and is nothing similar to a *node that is not configured as a compute node in a grid, in response to grid configuration information sent by a discovered master node in the grid, self-configuring as a compute node in the grid in accordance with the grid configuration information.* The cited references, alone or in combination, do not teach or suggest *wherein the node is further configured to, in response to said grid configuration information, self-configure as a compute node in the grid in accordance with the grid configuration information.*

Combining IEEE’s method of discovery with Doyle’s system **would not result in what is recited in claim 14 of the instant application.** For example, Doyle’s client computers, which include a “client control program,” are clearly **pre-configured** as “compute nodes” in Doyle’s grid. The Doyle reference does not teach or suggest that Doyle’s “client computers 11” have any need for discovery of the master computer. Doyle’s pre-configured “client computers 11”, from Doyle’s description, are already aware of the master computer, and thus Doyle’s system would have no need for discovery as disclosed in IEEE. **Furthermore, even if the Doyle and IEEE references were properly combinable, the combination would not produce anything like what is recited in claim 14 of the present application.**

Applicants also assert that one of ordinary skill in the art would have had no reason to combine the references in the manner proposed by the Examiner. For example, Doyle's client computers, which include a "client control program," are clearly **pre-configured** as "compute nodes" in Doyle's grid. The Doyle reference does not teach or suggest that Doyle's "client computers 11" have any need for discovery of the master computer. Doyle's pre-configured "client computers 11", from Doyle's description, are already aware of the master computer, and thus Doyle's system would have no need for discovery as disclosed in IEEE. Thus, the proposed combination makes no sense. (Regarding the Kampe reference, Kampe already discloses locating other cluster nodes, so there is no reason to combine Kampe with IEEE for discovery purposes.) Also, as shown above, the IEEE reference teaches away from master node discovery and does not pertain to the types of systems in Doyle and Kampe. Moreover, the Examiner's stated reason for combining the Doyle and IEEE references is merely conclusory. And in any case, as noted above, the cited references, alone or in combination, do not teach all of the limitations found in claim 14 of the instant application.

Thus, for at least the reasons presented above, the rejection of claim 14 is not supported by the cited art and removal thereof is respectfully requested.

Regarding claim 18, Applicants traverse the rejection of this claim for at least the reasons given above regarding claim 14.

Regarding claim 19, Applicants traverse the rejection of this claim for at least the reasons given above regarding claim 1.

Regarding claim 27, Applicants traverse the rejection of this claim for at least the reasons given above regarding claim 1.

The Examiner rejected claims 8, 12, 16, 25 and 33 under 35 U.S.C. § 103(a) as being unpatentable over Doyle, IEEE and Kampe, and further in view of Sun Cluster Grid Architecture (Sun Cluster Grid Architecture - a Technical White Paper Describing

the Foundation of Sun Grid Computing, published by Sun Microsystems on May 2002), and claims 9, 13, 17, 26 and 34 as being unpatentable over Doyle, IEEE and Kampe, and further in view of JXTA Chapter 1. Since the rejections have been shown to be unsupported for the independent claims, further discussion of these § 103(a) rejections is not necessary at this time.

Applicants also assert that numerous ones of the dependent claims recite further distinctions over the cited art. However, since the rejections have been shown to be unsupported for the independent claims, a further discussion of the dependent claims is not necessary at this time.

CONCLUSION

Applicants submit the application is in condition for allowance, and an early notice to that effect is requested.

If any fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5681-69600/RCK.

Respectfully submitted,

/Robert C. Kowert/

Robert C. Kowert, Reg. #39,255
Attorney for Applicant(s)

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8850

Date: November 26, 2008